




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 11/14	A1	(11) International Publication Number: WO 99/66401 (43) International Publication Date: 23 December 1999 (23.12.99)
(21) International Application Number: PCT/US99/13784 (22) International Filing Date: 17 June 1999 (17.06.99) (30) Priority Data: 09/099,844 19 June 1998 (19.06.98) US (71) Applicant: NETWORK APPLIANCE, INC. [US/US]; 2770 San Tomas Expressway, Santa Clara, CA 95051 (US). (72) Inventor: MUHLESTEIN, Mark; 15105 Venetian Way, Morgan Hill, CA 95037 (US). (74) Agent: SWERNOFSKY, Steven, A.; Swernofsky Law Group, P.O. Box 390013, Mountain View, CA 94039-0013 (US).		(81) Designated States: CA, CN, JP, KR, European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE). Published <i>With international search report.</i> <i>Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.</i>

(54) Title: BACKUP AND RESTORE FOR HETEROGENEOUS FILE SERVER ENVIRONMENT



(57) Abstract

The invention provides a file server that provides for backup and restore of files in a heterogeneous file server environment. Metadata associated with each file by each model or protocol is preserved across backup and restore operations. The file server performs at least three tasks as part of the backup and restore operations. The file server notes all file attributes associated with each file in either model. These file attributes can include Unix perms, ownership and timestamp information, and link information. On backup, the file server records those file attributes in an analogous record using a different model. The file server records each element of Unix metadata as an NT Extended Attribute for use by an NT backup element. On restore, the file server restores those file attributes from the analogous record. The NT Extended Attributes provided by an NT restore element are translated into Unix metadata. The file server provides an element for translating between Unix metadata and NT Extended Attributes transparently to the NT backup and restore elements, so that NT backup and restore elements that were created without reference to the heterogeneous file server environment can be used.

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece			TR	Turkey
BG	Bulgaria	HU	Hungary	ML	Mali	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MN	Mongolia	UA	Ukraine
BR	Brazil	IL	Israel	MR	Mauritania	UG	Uganda
BY	Belarus	IS	Iceland	MW	Malawi	US	United States of America
CA	Canada	IT	Italy	MX	Mexico	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NE	Niger	VN	Viet Nam
CG	Congo	KE	Kenya	NL	Netherlands	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NO	Norway	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	NZ	New Zealand		
CM	Cameroon			PL	Poland		
CN	China	KR	Republic of Korea	PT	Portugal		
CU	Cuba	KZ	Kazakhstan	RO	Romania		
CZ	Czech Republic	LC	Saint Lucia	RU	Russian Federation		
DE	Germany	LI	Liechtenstein	SD	Sudan		
DK	Denmark	LK	Sri Lanka	SE	Sweden		
EE	Estonia	LR	Liberia	SG	Singapore		

Title of the Invention

Backup and Restore for Heterogeneous File Server Environment

5

Background of the Invention1. *Field of the invention*

The invention relates to storage systems.

10

2. *Related Art*

15

A file servers is a computer storage systems that is used to record and retrieve data, in response to requests made by client devices; these client devices are typically coupled to one or more file servers using a communication network. The file server receives requests formatted in a file server protocol and following a data storage and retrieval model that is suited to that protocol. (In many circumstances the file server protocol and the data storage and retrieval model are nearly inseparable.)

20

Because different client devices often use differing data storage and retrieval models, or differing file server protocols, it is desirable for the file server to provide services using a plurality of such models and protocols. One such model is the Unix file system, with the associated NFS (Network File System) protocol. Other file systems include the Windows NT file system, with the associated CIFS (Common Internet File System) protocol.

25

One problem that has arisen in the art is that certain functions the file server should advantageously perform, such as backup and restore, differ substantially between differing models and protocols. It would be advantageous for the file server to perform these functions for files using either model, while preserving all the information associated with the file in either model. In particular, this poses a problem because some of the metadata functions for files using either model, while preserving all the information associated with the file in either model. In particular, this poses a problem because some of the metadata associated with files in the Unix file system is not represented in the Windows NT file system.

30

Accordingly, it would be advantageous to provide a file server that provides for backup and restore of files in a heterogeneous file server environment, in which metadata associated with each file by each model or protocol is preserved across backup and restore operations. This advantage is achieved in an embodiment of the invention in which the file server (a) notes all file attributes associated with each file; (b) records those file attributes in an analogous record on backup using a different model; and (c) restores those file attributes from the analogous record on a restore operation.

Summary of the Invention

The invention provides a file server that provides for backup and restore of files in a heterogeneous file server environment. Metadata associated with each file by each model or protocol is preserved across backup and restore operations. The file server performs at least three tasks as part of the backup and restore operations. The file server notes all file attributes associated with each file in either model. In a preferred embodiment, these file attributes can include Unix perms, ownership and timestamp information, and link information. On backup, the file server records those file attributes in an analogous record using a different model. In a preferred embodiment, the file server records each element of Unix metadata as an NT Extended Attribute for use by an NT backup element. On restore, the file server restores those file attributes from the analogous record. In a preferred embodiment, the NT Extended Attributes provided by an NT restore element are translated into Unix metadata.

In a preferred embodiment, the file server provides an element for translating between Unix metadata and NT Extended Attributes transparently to the NT backup and restore elements, so that NT backup and restore elements that were created without reference to the heterogeneous file server environment can be used.

Brief Description of the Drawings

Figure 1 shows a block diagram of a file server system providing backup and restore in a heterogeneous environment.

Figure 2 shows a block diagram of a file server system providing backup and restore capability in a heterogeneous environment.

Figure 3 shows a process flow diagram of a method of operation for a file server system providing backup and restore in a heterogeneous environment.

5

Detailed Description of the Preferred Embodiment

In the following description, a preferred embodiment of the invention is described with regard to preferred process steps and data structures. However, those skilled in the art would recognize, after perusal of this application, that embodiments of the invention may be
10 implemented using one or more general purpose processors (or special purpose processors adapted to the particular process steps and data structures) operating under program control, and that implementation of the preferred process steps and data structures described herein using such equipment would not require undue experimentation or further invention.

15

Figure 1 shows a block diagram of a file server system in a heterogeneous environment.

A system 100 includes a communication network 110, a set of NT workstations 120, a set of Unix workstations 120, an NT application server 130, a Unix application server
20 130, and a multi-protocol file server 140.

25

The communication network 110 can include a LAN (local area network) or other communication structure or technique for coupling the workstations 120, the application servers, and the file server 140.

The workstations 120 include devices that take the role of client device in a client-server model of application use and file storage. NT workstations 120 can be PC devices using the "Windows 95" operating system, the NT operating system, or another variant of the "Windows: operating system. Unix workstations 120 can be devices using any variant of the
30 Unix operating system, or other operating systems bearing sufficient similarity to Unix to interoperate herewith.

The application servers 130 can be general-purpose or special-purpose devices that take the role server device in a client-server model of application use.

The file server 140 includes a processor, program and data memory, and mass storage for recording and retrieving data in a file system. In a preferred embodiment, the file server 140 is disposed for receiving and responding to file server requests using the communication network. The file server requests can use either the CIFS file server protocol (and accordingly, an NT file system model) or the NFS file server protocol (and accordingly, a Unix file system model). In alternative embodiments, other file server protocols and file system models may be supported.

10 *Backup and Restore Capability*

Figure 2 shows a block diagram of a file server system providing backup and restore capability in a heterogeneous environment.

15 A system 200 includes a set of file servers 140, a set of backup storage devices 210, and a backup control device 220.

In a preferred embodiment, the file servers 140 are as described with regard to figure 1. In alternative embodiments, the file servers 140 may include single protocol file servers 140 or multi-protocol file server 140 that support only some of the protocols and file system models described herein.

The backup storage devices 210 are coupled to the file servers 140 using the communication network 110, and are disposed for recording and retrieving backup information from the file servers 140.

In a preferred embodiment, the backup storage devices 210 can include one or more magnetic, magneto-optical, or optical disks in a "jukebox" configuration.

30 In a first set of alternative embodiments, the backup storage devices 210 can include one or more tape drives or other mass storage devices. In a second set of alternative embodiments, the backup storage devices 210 can include a set of file servers 140 disposed for acting in a backup and restore capacity.

The backup control device 220 is coupled to at least some of the file servers 140 using the communication network 110, and is disposed for controlling those file servers 140 for backup and restore operations.

5 *Method of Operation*

Figure 3 shows a process flow diagram of a method of operation for a file server system providing backup and restore in a heterogeneous environment.

10 A method 300 is performed by the system 100 operating in conjunction, including a set of file servers 140, the backup storage device 210, and the backup control device 220.

File Metadata

15 The method 300 operates to preserve file metadata across backup and restore operations.

File metadata includes information pertinent to each file, not already included in the file data (the file data is typically the contents of the file). For example, file metadata can
20 include one or more of the following:

one or more names for the file;

access control information for the file, such as (a) Unix "permissions" bits, and related
25 bits such as the "use-owner-permissions" bit, or (b) Unix group identifiers or owner identifiers for the file;

linking information for the file, such as (a) whether the file is a symbolic link or a Unix
"hard" link, or (b) the target file or Unix inode for the link; or

30 other possible metadata for the file, such as timestamps for when the file was last read from or written to.

The scope and spirit of the invention are not limited to these specific examples. Those skilled in the art will recognize, after perusal of this application, that other and further information about a file may be considered file metadata in embodiments of the invention.

5 In a preferred embodiment, to perform metadata-conversion, the file server 140 performs the following steps:

 At a flow point 350, the file server 140 is ready to perform a metadata-conversion operation for a selected file.

10 At a step 351, the file server 140 notes each item of metadata associated with the selected file.

 At a step 352, the file server 140 records each item of metadata in an analogous
15 record using the NT file system model. In a preferred embodiment, the file server 140 records each item of metadata as an NT Extended Attribute for use by an NT backup element.

 In alternative embodiments, the file server 140 may record each item of metadata using other aspects of the NT file system model. For example, the file server 140 may use one
20 or more of the following techniques:

 The file server 140 may create one or more associated files with special names or including the associated metadata;

25 The file server 140 may insert a header or footer in the file to record the items of metadata; and

 The file server 140 may create a new NT data stream specific to the purpose of recording and retrieving the items of metadata.

30 At a flow point 360, the metadata-conversion operation is performed as part of the backup function.

Backup Function

At a flow point 310, the system 100 is ready to perform a backup function.

5 At a step 311, the backup control device 220 directs one of the file servers 140 to perform a backup operation.

 At a step 312, the file server 140 determines the file system model for files in its file system (or subset of its file system) for which the backup control device 220 directed it to perform the backup operation. If the file system model for any particular file is the NT file system model, no special processing of file metadata occurs for that file. If the file system model for any particular file is the Unix file system model, the file server 140 performs the step 313, at which it performs the metadata-conversion operation described with regard to flow point 350.

15 At a step 313, the file server 140 performs the metadata-conversion operation described with regard to flow points 350 through 360.

 In a preferred embodiment, the metadata-conversion operation is transparent to NT backup and restore elements. This allows NT backup and restore elements that were created without reference to the heterogeneous file server environment to be used without substantial change.

 The metadata-conversion operation is somewhat different if the file from the Unix file system is actually a Unix hard link. Windows NT backup and restore utilities are not programmed to properly handle multiple file names referring to the same file. Accordingly, for Unix hard links, the metadata-conversion operation saves sufficient information to allow NT backup utilities to backup the file, and to allow the Unix hard link to be reconstructed after operation of NT restore utilities.

30 In a preferred embodiment, when the file is a Unix hard link, the metadata-conversion operation tracks references to those files (by reference to the particular Unix inode), and converts all references except the first to hard links. The metadata-conversion operation adds NT Extended Attributes sufficient to identify the hard links upon restore.

At a flow point 320, the backup function is complete.

Restore Function

5 At a flow point 330, the system 100 is ready to perform a restore function.

At a step 331, the file server 140 reverses the backup operation, including the metadata-conversion operation performed in the step 313.

10 Those skilled in the art will recognize, after perusal of this application, how to reverse the metadata-conversion operation described with regard to flow points 350 through 360. The file server 140 can determine those items of metadata recorded during the metadata-conversion operation, retrieve those items of metadata for each selected file, and record those items of metadata for the file using the Unix file model.

15 At a flow point 340, the restore function is complete.

Other Utilities

20 It will be clear to those skilled in the art, after perusal of this application, that the invention has wide applicability, and is not limited to backup and restore functions.

For example, the invention may be used in conjunction with other programs that are used to copy files or directories, in which it is advantageous to preserve metadata associated with those files or directories. The invention is particularly advantageous when it is desired to use software utilities that do not recognize the metadata to be preserved.

Generality of the Invention

30 The invention has wide applicability to other operations in a multi-protocol file server. For example, the invention may be used with one or more of the following other techniques:

The invention may be used to record items of metadata from the NT file system model in the Unix file system model, and to retrieve those items of metadata for restoring files using the NT file system model. This would allow use of Unix backup and restore elements in addition to or instead of NT backup and restore elements.

5

The invention may be used in conjunction with inventions described in the following application:

10

Application Serial No. 09/037,652, filed March 10, 1998, in the name of the inventor Steven Kleinman, titled "Highly Available File Servers," attorney docket number NAP-012.

This application is hereby incorporated by reference as if fully set forth herein.

15

The invention may be used to backup and restore system records and other items of metadata not specifically associated with any particular file. These system records may include file system structure information, process control information, and other information relevant to operation of the file server 140.

20 *Alternative Embodiments*

Although preferred embodiments are disclosed herein, many variations are possible which remain within the concept, scope, and spirit of the invention, and these variations would become clear to those skilled in the art after perusal of this application.

Claims

1. A file server,
said file server disposed for recording and retrieving each one of a set of files in
5 one of a plurality of data storage and retrieval models;
each one of said data storage and retrieval models disposed for including
metadata with at least some of said files;
said file server including a backup element disposed for preserving said metadata
for retrieval associated with said files even when said files are recorded by said backup element
10 in one of said data storage and retrieval models other than associated with said files.
2. A file server as in claim 1, wherein said backup element records said
metadata as an NT Extended Attribute for use by an NT backup element.
- 15 3. A file server as in claim 1, wherein said backup element records said
metadata in an analogous record using a data storage and retrieval model other than associated
with said files.
4. A file server as in claim 1, wherein said data storage and retrieval models
20 include Unix/NFS and NT/CIFS.
5. A file server as in claim 1, wherein said file server includes an element
disposed for translating between Unix metadata and NT Extended Attributes transparently to a
set of NT backup and restore elements.
25
6. A file server as in claim 1, wherein said metadata includes at least one of

Unix permission bits;
ownership or group identifiers;
30 timestamp information; and
file link information.
7. A file server as in claim 1, wherein said metadata includes file link
information; and said backup element distinguishes between hard links

and symbolic links.

8. A file server as in claim 1,

Including a restore element disposed for recovering said metadata associated with said files even when said files are recorded by said backup element in one of said data storage and retrieval models other than associated with said files;

wherein said restore element retrieves said metadata from said analogous record and records said metadata using a data storage and retrieval model associated with said files.

9. A file server as in claim 8, wherein said metadata includes file link information; and

said restore element maintains connectivity specified by hard links in said set of files.

10. A method of operating a file server, said method including steps for recording and retrieving each one of a set of files in one of a plurality of data storage and retrieval models;

including metadata with at least some of said files in each one of said data storage and retrieval models;

performing a backup operation, including steps for preserving said metadata for retrieval associated with said files even when said files are recorded by said backup element in one of said data storage and retrieval models other than associated with said files.

11. A method as in claim 10, including steps for maintaining connectivity specified by hard links in said set of files.

12. A method as in claim 10, including steps for recovering said metadata associated with said files even when said files are recorded by said backup element in one of said data storage and retrieval models other than associated with said files;

wherein said steps for recovering include steps for retrieving said metadata from said analogous record, and steps for recording said metadata using a data storage and retrieval model associated with said files.

13. A method as in claim 10, including steps for translating between Unix metadata and NT Extended Attributes transparently to a set of NT backup and restore elements.

5 14. A method as in claim 10, wherein said steps for preserving include steps for recording said metadata as an NT Extended Attribute for use by an NT backup element.

15. A method as in claim 10, wherein said steps for preserving include steps for recording said metadata in an analogous record using a data storage and retrieval model other
10 than associated with said files.

1/2




FIG. 1




FIG. 2

2/2

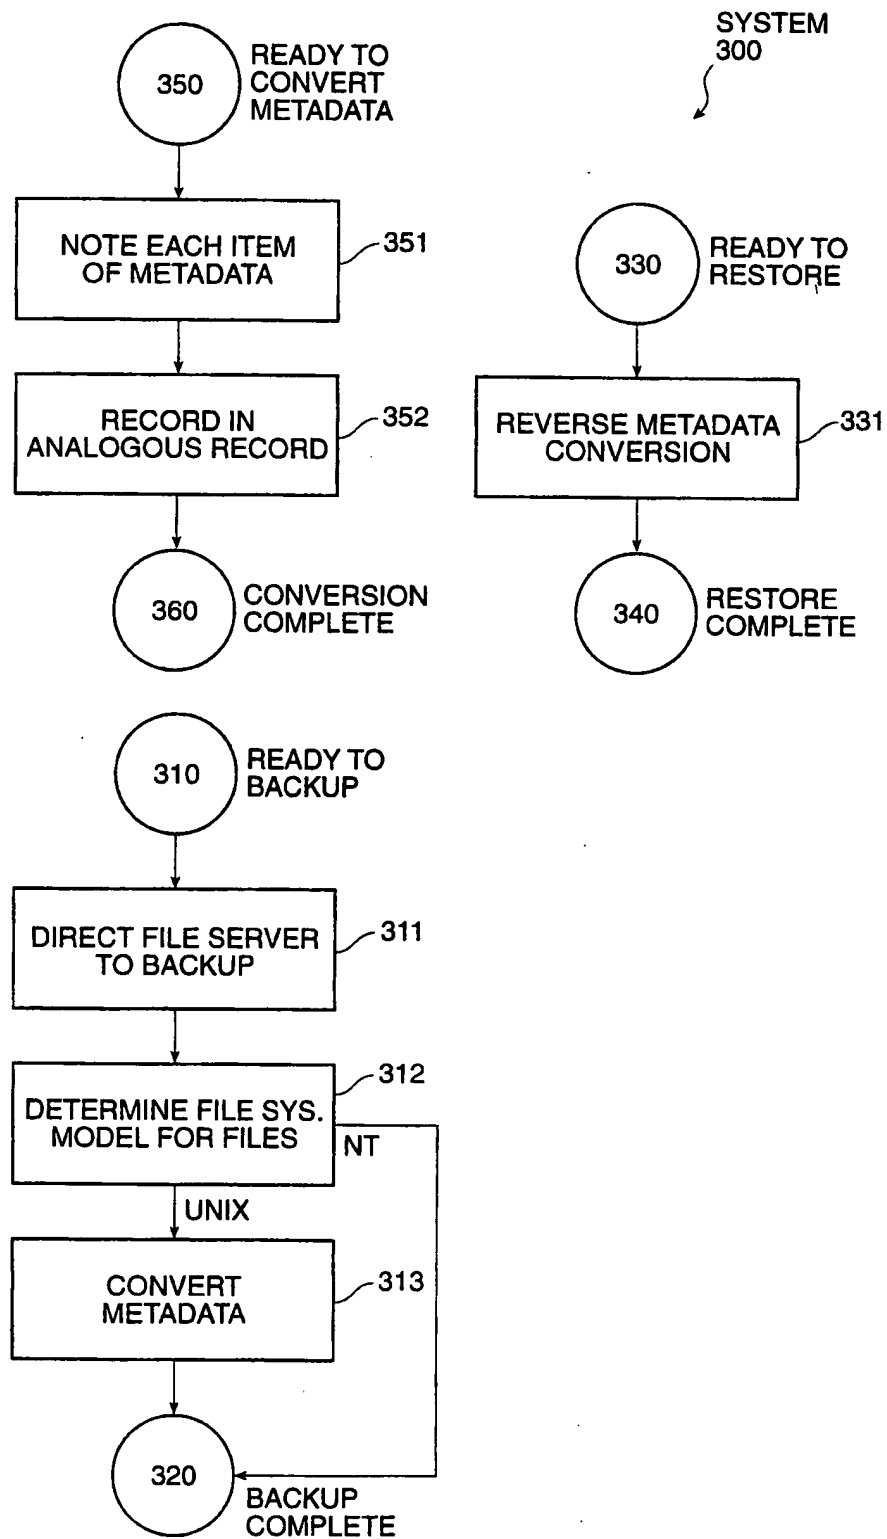


FIG. 3

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 99/13784

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F11/14

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	"Mapping the VM text files to the AIX text files" IBM TECHNICAL DISCLOSURE BULLETIN., vol. 33, no. 2, July 1990 (1990-07), page 341 XP000123641 IBM CORP. NEW YORK., US ISSN: 0018-8689 the whole document	1-15
A	"Migrated Data Backup Utility" IBM TECHNICAL DISCLOSURE BULLETIN., vol. 37, no. 06B, June 1994 (1994-06), pages 505-507, XP000456079 IBM CORP. NEW YORK., US ISSN: 0018-8689	1-15

☐ Further documents are listed in the continuation of box C.☐ Patent family members are listed in annex.

* Special categories of cited documents :

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier document but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

"&" document member of the same patent family

Date of the actual completion of the international search

11 October 1999

Date of mailing of the international search report

18/10/1999

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Corremans, G